




AIF-C01 Domain 1: Fundamentals of AI and ML

AIF-C01 exam guide Domain 1 (20 %), as of 2026 | 1.1 AI terms and learning types, 1.2 practical use cases and when not to use ML, AWS task AI services, 1.3 the ML lifecycle, SageMaker AI components, MLOps, model and business metrics

AIF-C01 Domain 1 ML lifecycle SageMaker AI Metrics


1.1 Basic AI concepts

AI contains ML contains deep learning contains generative AI



Learning type	Data	Tasks	Example
Supervised	labelled	classification (category), regression (number)	spam or not; house price
Unsupervised	unlabelled	clustering, dimensionality reduction, anomaly detection	customer segments; unusual transactions
Semi-supervised	few labels, much unlabelled	label propagation	medical images with few annotations
Self-supervised	raw data creates its own labels	pre-training (predict the next word)	how FMs learn
Reinforcement	rewards from an environment	sequential decisions	robot control, DeepRacer, RLHF

A small classifier network



Data types — Labelled vs unlabelled; structured (tables, time series) vs unstructured (text, images, audio, video).

Inference types — Real-time (one request, low latency), batch (large datasets offline), asynchronous (large payloads, queued), serverless (spiky traffic, scales to zero).

Computer vision and NLP — CV extracts meaning from images and video; NLP from human language (sentiment, entities, translation).


1.2 Practical use cases and AWS services

ML adds value when	Do not use ML when
patterns are too complex for rules (fraud, images, language)	a deterministic rule gives the exact answer (tax formula)
decisions repeat at scale (millions of predictions)	cost of building exceeds the benefit
data exists and keeps arriving	there is no (or biased) data
a probabilistic answer is acceptable	every decision must be fully explainable by law

Need	AWS managed AI service
Detect objects, faces, unsafe content in images/video	Amazon Rekognition
Extract text, forms, tables from documents	Amazon Textract
Sentiment, entities, key phrases, PII in text	Amazon Comprehend
Speech to text (call transcripts, subtitles)	Amazon Transcribe
Text to lifelike speech	Amazon Polly
Translate text between languages	Amazon Translate

Chatbots and voice bots with intents	Amazon Lex
Recommendations and personalisation	Amazon Personalize
Enterprise search over documents	Amazon Kendra
Build, train, deploy custom models	Amazon SageMaker AI

1.3 The ML development lifecycle



SageMaker AI component	Lifecycle step
Ground Truth	label data with humans (and active learning)
Data Wrangler / Canvas data prep	explore, clean and transform data with little code
Feature Store	store and share features (online and offline)
JumpStart	pre-trained and foundation models, solution templates
Training jobs, automatic model tuning	train at scale, search hyperparameters
Clarify	bias detection and explainability (feature attributions)
Model Registry, Pipelines	version, approve and automate (MLOps)
Endpoints (real-time, serverless, async), batch transform	inference
Model Monitor	data and model quality drift alerts
Canvas	no-code ML for business analysts

Evaluate a churn classifier (1,000 customers)

Actual \ Predicted	Predicted		Total
	Churn	Stay	
Churn	60 TP	40 FN	100
Stay	30 FP	870 TN	900
Total	90	910	1000
Precision	0.67	0.96	
Recall	0.60	0.97	
F1	0.63	0.96	

Accuracy 93% · n = 1000

Regression metrics by hand

Y A model predicts delivery time in days. On 4 orders the errors (predicted - actual) are 2, -1, 3 and 0. Compute MAE and RMSE.
Given: $e = 2, -1, 3, 0$ days ($n = 4$); MAE = mean of $|e|$; RMSE = square root of the mean of e^2

Y Step 1: MAE?

$(2 + 1 + 3 + 0) / 4 = 1.5$ days
average size of the error **T**

Y Step 2: Squared errors?

$4, 1, 9, 0 \rightarrow \text{mean } 14 / 4 = 3.5$
squaring punishes big errors **T**

Y Step 3: RMSE?

$\sqrt{3.5} = 1.87$ days
back to the unit of the target **T**

Y Step 4: Why RMSE > MAE?

$1.87 > 1.5$
one large error (3 days) weighs more in RMSE **T**


MAE = 1.5 days, RMSE ≈ 1.87 days **T**
Check: Use RMSE when big misses are costly; R^2 tells the share of variance explained.

- **Model metrics** classification: accuracy, precision, recall, F1, AUC-ROC; regression: MAE, RMSE, R^2
- **Business metrics** cost per user or prediction, ROI, revenue uplift, churn reduced, customer feedback, development cost

- **MLOps concepts** experiment tracking, repeatable pipelines, scalable systems, managing technical debt, production readiness, monitoring and retraining
- **Model sources** use a pre-trained or open-source model (JumpStart, Bedrock) or train a custom one; serve through a managed API or a self-hosted endpoint

🚩
 Recap and self-test

Statement	True	False
1. Clustering is supervised learning → <i>unsupervised: no labels</i>	☐	☒
2. Predicting a price is regression → <i>a continuous number</i>	☒	☐
3. Batch inference suits a nightly scoring of 10 million records → <i>no low-latency need</i>	☒	☐
4. Accuracy is the best metric for rare fraud → <i>recall, precision, F1 matter on imbalanced data</i>	☐	☒
5. Model Monitor detects data drift on endpoints → <i>compares live data to a baseline</i>	☒	☐
6. Hyperparameters are learned from the data during training → <i>they are set before training; weights are learned</i>	☐	☒



Q1
 Group customers by behaviour with no labels: which learning type?

 A
 Unsupervised learning (clustering).

Q2
 Which service labels training data with human workers?

 A
 SageMaker Ground Truth.

Q3
 Which SageMaker AI capability detects bias and explains predictions?

 A
 SageMaker Clarify.

Q4
 Inference type for large payloads taking minutes, scale to zero?

 A
 Asynchronous inference.

Q5
 Precision from TP = 60, FP = 30?

 A
 $60 / (60 + 30) = 0.67$.

Q6
 Recall from TP = 60, FN = 40?

 A
 $60 / (60 + 40) = 0.60$.

Q7
 Which service turns call audio into text?

 A
 Amazon Transcribe.

Q8
 No-code ML tool for business analysts?

 A
 SageMaker Canvas.

Q9
 Name a business metric for an ML project.

 A
 ROI, cost per prediction, revenue uplift, churn reduction or customer satisfaction.